

Teadmiste formaliseerimine

aka

Knowledge Representation
(and Reasoning)

KR & R

Intro

Course page

http://lambda.ee/wiki/Teadmiste_formaliseerimine

Course theme

From SQL to natural language

Course theme

From SQL to natural language

- Structured facts: SQL
- Unstructured knowledge: natural language
- *Focus of the course: semi-structured knowledge*

Course theme

From SQL to natural language

- Structured facts: SQL
- Unstructured knowledge: natural language
- *Focus of the course: semi-structured knowledge*

Capturing meaning in natural language sentences:

understanding what is said, answering questions, explaining answers, calculating plans of action.

Journey: main steps of the course

SQL and logic

SQL and triples / RDF / JSON-LD

Simple rules like in RDFa

Large semi-structured knowledge bases

Machine learning in NLP: vector representation of words

Machine learning in NLP: large language models (BERT, GPT)

Reasoning with uncertain knowledge / uncertain rules

Semantic parsing

Hybrid question answering: machine learning + rule-based reasoning

Wikipedia

In artificial intelligence (AI), [commonsense reasoning](#) is a human-like ability to make presumptions about the type and essence of ordinary situations humans encounter every day.

These assumptions include judgments about the nature of physical objects, taxonomic properties, and peoples' intentions.


Wikipedia AGI

Artificial general intelligence (AGI) is the ability of an intelligent agent to understand or learn any intellectual task that a human being can. It is a primary goal of some artificial intelligence research and a common topic in science fiction and futures studies.

AGI can also be referred to as **strong AI**, **full AI**, or **general intelligent action**, although some academic sources reserve the term "strong AI" for computer programs that experience sentience or consciousness.

What is pure classical KR & R?

- A classic old-school A.I. subfield *for actual human-like thinking*
- Used to set lofty goals extremely hard to achieve (and did not)
- Famous A.I. winters:



Approach to thinking in KR&R

Symbolic A.I:

- Write down (human?) knowledge using rules and facts.
- Answer questions and solve problems by deriving new knowledge using rules.

Next, small examples using an automated reasoner for first order logic:
run these yourself in <http://logictools.org>

spacity(kuressaare).

railroad(tallinn,tapa).

railroad(tapa,tartu).

highway(tallinn,virtsu).

sealane(virtsu,kuivastu).

highway(kuressaare,kuivastu).

railroad(X,Y) => railroad(Y,X).

highway(X,Y) => highway(Y,X).

sealane(X,Y) => sealane(Y,X).

railroad(X,Y) => easytravel(X,Y,use(train,X,Y)).

highway(X,Y) => easytravel(X,Y,use(bus,X,Y)).

sealane(X,Y) => easytravel(X,Y,use(ship,X,Y)).

easytravel(X,Y,P1) & easytravel(Y,Z,P2) => easytravel(X,Z,combine(P1,P2)).

easytravel(tartu,X,Y) & spacity(X) => \$ans(X,Y).

- Replace *spacity(kuressaare)* with

city(kuressaare).

has(kuressaare,coast).

hasnrof(kuressaare,spa)=4.

population(kuressaare)=10000.

$\$less(population(X),20000) \Rightarrow smallcity(X).$

$(city(X) \ \& \ has(X,coast) \ \& \ \$less(3,hasnrof(X,spa)) \ \& \ smallcity(X)) \Rightarrow spacity(X).$

...

Note: it can be pretty hard to find solutions

- Finding a solution is semidecidable: no guarantee that we can show that a solution does *not* exist.
- Even small problems can be hard. A solvable tiny example from group theory, ca 57 steps in a proof (try it out on logictools.org):

$\text{multiply}(A, \text{inverse}(\text{multiply}(B, \text{multiply}(C, \text{multiply}(\text{multiply}(\text{inverse}(C), \text{inverse}(\text{multiply}(D, B))), A)))) = D.$

$\text{multiply}(\text{inverse}(a1), a1) \neq \text{multiply}(\text{inverse}(b1), b1).$

What happens with this approach?

- Small toy problems are easy to present and solve.
- Presenting and solving large complex problems for common-sense reasoning has turned out to be extremely hard.

Why hard

- Not always easy to find solutions.
- Extremely hard to build up a knowledge base: so many rules needed
- Standard logic is too rigid and not expressive enough:
 - How to represent and perform reasoning with confidences / probabilities?
 - How to handle default rules (i.e. rules with exceptions)?
 - How to handle time, space and context?
 - How to handle ongoing changes in the represented world?
 - How to handle reasoning by analogy?
 - How to connect to natural language?
 - ...
- **None** of these questions have good answers or well-working systems so far

A narrow subfield of KR &R: automated reasoning

- AKA “Automated Theorem Proving”
- Focus on developing algorithms and systems for automatically solving problems written in logic.
- Yearly competitions <http://www.tptp.org/CASC/> and the main conference CADE <http://www.cadeinc.org/conferences>
- Focus more on mathematical kinds of problems, not so much common sense or very large knowledge bases.
- In other words, top systems work best with relatively few rules and deep complex proofs.
- Not too many high-performance systems (like, my system Gandalf was world leading at the end of 20th century for many years, and the current GKC system is one of the top systems nowadays)

Classic old-style approach

Three components needed:

- English to logic parser
- Logical reasoner
- Knowledge bases used by the reasoner

They all failed miserably!

Famous case study: CYC



Wikipedia

- **Cyc** is a long-term artificial intelligence project that aims to assemble a comprehensive ontology and knowledge base that spans the basic concepts and rules about how the world works.
- Douglas Lenat began the project in July 1984 ... since January 1995, has been under active development by the **Cycorp** company.
- The Cyc project has been described as "[one of the most controversial endeavors of the artificial intelligence history](#)". Catherine Havasi, CEO of Luminoso, says that Cyc is the predecessor project to IBM's Watson. Machine-learning scientist Pedro Domingos refers to the project as a "catastrophic failure".

A mostly failed limited approach: Semantic Web

- Idea: make it easy to publish, harvest, integrate and use structured knowledge on the web.
- Proposals worked out: RDF data model and OWL rule-like language
- RDF main idea: represent all data in a single three-column table of <object id> <property> <value> triplets like this:

client_1 name „John Brown“

client_1 balance 200

What happened with the Semantic Web?

Negative:

- Mostly academic research leading to over-complicating even simple things.
- People were not really keen to publish data or rules on the web.

Positive:

- Attempt to make it a bit simpler in a closely related area: linked data
- JSON-LD language giving RDF interpretation to JSON
- <http://schema.org> repository of suggested type/property names

Positive case studies of KR & R

- Google's Knowledge Graph used e.g. for infoboxes and Google Assistant
- Facebook Knowledge Graph
- Apple Siri knowledge base
- IBM Watson question-answering system from 2011

You may want to read a big overview of knowledge graphs:

<https://arxiv.org/pdf/2003.02320.pdf>

IBM Watson winning Jeopardy! in 2011



Example questions from Jeopardy with Watson

- It's just a bloody nose! You don't have this hereditary disorder once endemic to European royalty : *Haemophilia*
- You just need a nap! You don't have this sleep disorder that can make sufferers nod off while standing up : *Narcolepsy*
- Heitor Villa-Lobos dedicated his "12 Etudes" for this instrument to Andres Segovia : *Guitar*
- Paganini's "24 Capricci" set the standard for etudes for this instrument : *Violin*
- Rembrandt's Biblical Scene "Storm on the Sea of" this was stolen from a Boston museum in 1990 : *Galilee*

End-to-end deep learning for Q&A

Main systems used by people for examples & research:

- Derivates of google [BERT](#): RoBERTa, DistilBERT, DeBERTa, ...
- OpenAI [GPT-3, ChatGPT](#)
- Newer Facebook and Google systems: not public yet

Exploring GPT-3: some famous benchmarks

Is it easier for human's to survive in an ocean, cave, town or alone?

It is easier for human's to survive in a town.

Metal is a thermal conductor. Which of these would let the most heat travel through? A new pair of jeans, a steel spoon in a cafeteria, a cotton candy at a store or a calvin klein cotton hat?

A steel spoon in a cafeteria would let the most heat travel through.

The large ball crashed right through the table because it was made of steel. What was made of steel?

The large ball.

The large ball crashed right through the table because it was made of styrofoam. What was made of styrofoam?

The table was made of styrofoam.

Exploring GPT-3 and ChatGPT

John is a man or not a man?

GPT3: I cannot answer that question. ChatGPT: John is a man.

John has a car and does not have a car?

GPT3: John has a car. ChatGPT: ... contradiction

John is nicer than Mike. Mike is nicer than Eve. Who is nicer than John?

GPT3 and ChatGPT: Eve is nicer than John.

Exploring GPT-3 and ChatGPT:

Penguins are birds who do not fly. Birds fly. John is a penguin. John flies?

GPT3 and ChatGPT: False. Penguins are birds who do not fly.

John is a bird. Folkameers are birds who can not fly. Birds can fly. John is not a Folkameer. John can fly?

GPT3: No, John is not a Folkameer. ChatGPT: No

Penguins cannot fly. Most birds can fly. Most things cannot fly. Birds are things. John is a bird. Can John fly?

ChatGPT: No

Penguins cannot fly. Most birds can fly. Most things cannot fly. Birds are things. John is a bird. Why cannot John fly?

ChatGPT: John cannot fly because he is a penguin and penguins are birds that cannot fly.

Exploring GPT-3 and ChatGPT:

Two birds sat on a tree. Three more birds came to sit on the tree. How many birds are sitting on the tree now?

GPT3: There are five birds sitting on the tree now.

Ten birds sat on a tree. Seven more birds came to sit on the tree. How many birds are sitting on the tree now?

GPT3: There are now ten birds sitting on the tree. ChatGPT3: 17 birds.

1234 birds sat on a tree. 567 more birds came to sit on the tree. 91 birds left the tree. How many birds are sitting on the tree now?

ChatGPT: 1770 birds. (notice: the correct answer is 1710)

Exploring GPT-3 and ChatGPT:

A robot hand can lift one block at a time.

Birds can fly.

Penguins are birds which cannot fly.

John is a bird.

The robot hand cannot hold more than one block.

The robot hand can put a block on another block or a table.

A block can be on another block.

No more than one block can be on another block.

If a block X is on another block Y, the robot hand cannot lift the block Y.

There are three blocks on the table: a red block, a blue block, a yellow block.

Question: Can John fly?

GPT3 and ChatGPT: No

Question: Why cannot John fly?

GPT3 and ChatGPT: John cannot fly because he is a penguin.

Explainability?

Not much success explaining the reasoning behind the answers.

Possible argument against pure end-to-end learning:
brain has a genetic pre-determined structure

- Just observe a newborn calf.

https://www.youtube.com/watch?v=kUPX6p_btJ0

- Wikipedia: universal grammar (UG), in modern linguistics, is the theory of the [genetic component of the language faculty](#), usually credited to Noam Chomsky. The basic postulate of UG is that there are innate constraints on what the grammar of a possible human language could be. When linguistic stimuli are received in the course of language acquisition, children then adopt specific syntactic rules that conform to UG.

Can we fix end-to-end learning?

Nobody knows so far.

A lot of people put their hope on hybrid systems:

Combine learning with rule-based reasoning, arithmetic, SQL etc.

Most landmark A.I. systems are hybrid

... except for object recognition and story/picture generation

- Siri and Google Assistant
- Most automated driving systems
- AlphaGo and AlphaZero
- AlphaFold

Can we pre-build some necessary structure?

Like making the system internally know things like:

- There are objects, and they are mostly persistent.
- Objects are in space and time and follow some basic physics.
- Objects have properties.
- Objects have relations between each other.
- Relations have some basic logical properties.
- Some things are dangerous, some are desirable

...

Hybrid systems: ML + rules

Two approaches

- Inject rules into a machine learning model (extremely active field)
- Add machine learning to a rule-based system (only a few groups)

Inject rules into machine learning

- Immense number of possible ways to do that
- People are actively experimenting
- Typically we get small gains over pure end-to-end transformers

For example, try to calculate saliency and then count

Conversational Multi-Hop Reasoning with
Neural Commonsense Knowledge and
Symbolic Logic Rules

Forough Arabshahi
Facebook

Jennifer Lee
Facebook

Antoine Bosselut
EPFL

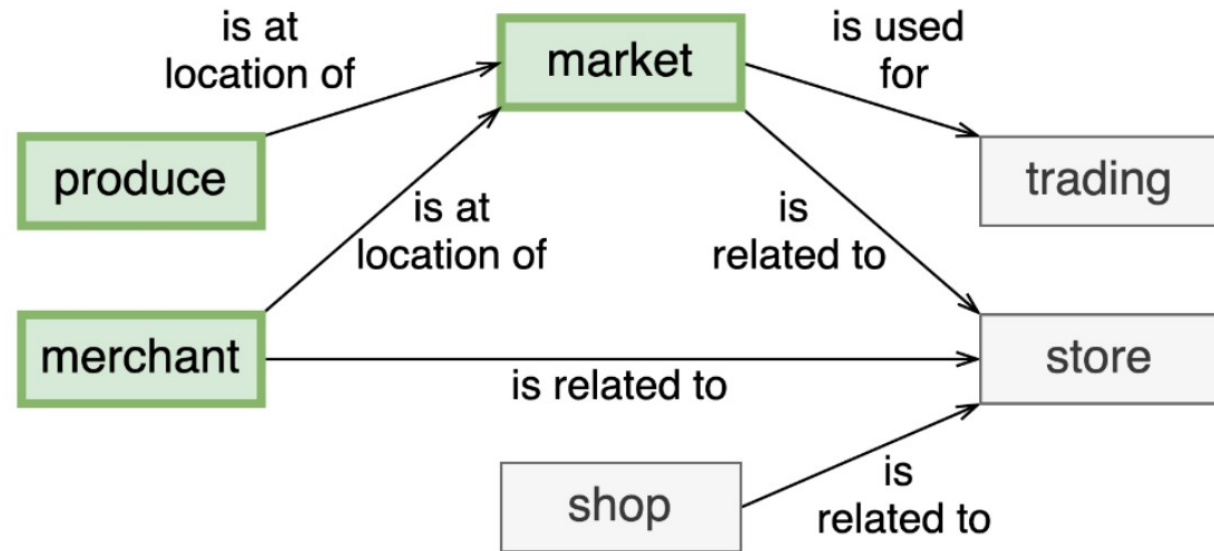
Yejin Choi
University of Washington

Tom Mitchell
Carnegie Mellon University

2021

What kind of store does a merchant have if they sell produce?

A. the sun **B. market** C. business D. mall E. shopping center



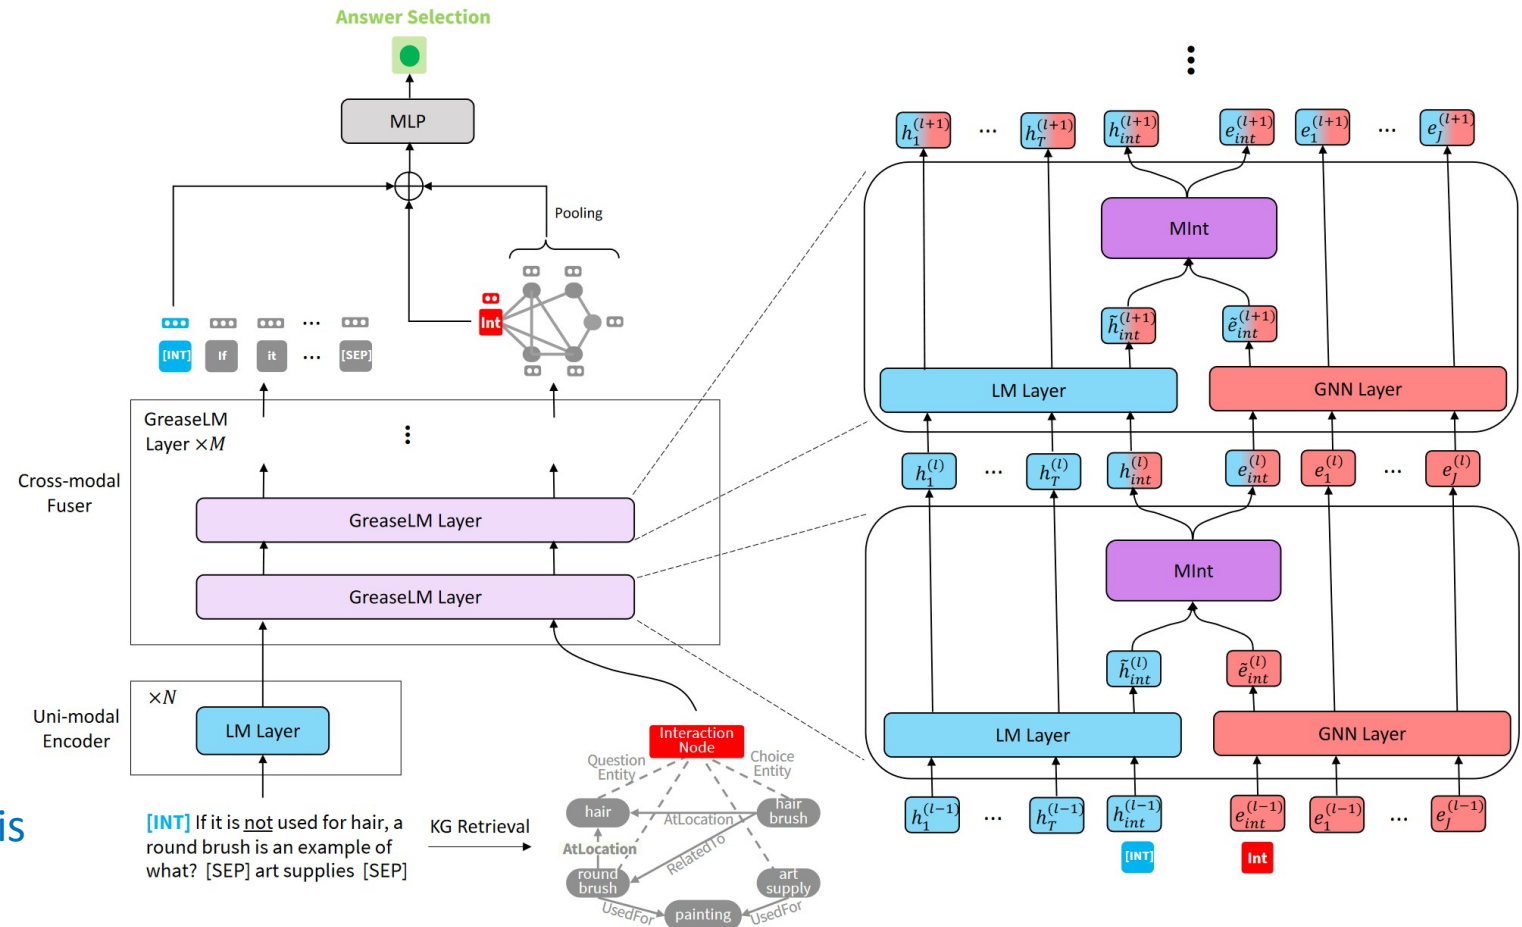
Example figure of one hybrid system

GREASELM: GRAPH REASONING ENHANCED LANGUAGE MODELS FOR QUESTION ANSWERING

Xikun Zhang, Antoine Bosselut, Michihiro Yasunaga, Hongyu Ren
Percy Liang, Christopher D. Manning, Jure Leskovec

Stanford University
2022

If it is not used for hair, a round brush is an example of what? (art supplies)




One big issue: good commonsense KBs?

Some of the most well-known:

- wordnet
- Dbpedia
- wikidata
- yago
- babelnet
- conceptnet
- atomic
- nell
- framenet
- cyc
- sumo
- TPTP
- schema.org

Typical KG contents: simple triplets like



Yet another active area: improve KBs

Conversational Multi-Hop Reasoning with
Neural Commonsense Knowledge and
Symbolic Logic Rules

Forough Arabshahi
Facebook


Jennifer Lee
Facebook

Antoine Bosselut
EPFL


Yejin Choi
University of Washington

Tom Mitchell
Carnegie Mellon University

2021



(a) ATOMIC beam results for two relations *Because I wanted* and *Before, I needed*



(b) ConceptNet beam results for two relations *requires action* and *is used for*

Case study: add machine learning to a rule-based system

Our group at Taltech: T. Tammet, P. Järv, D. Draheim, M. Verrev

Three components:

- A semantic parser: English to full first order logic
- A “softened” commonsense reasoner GK
- Improved versions of existing commonsense knowledge bases

Plus a minor thing: readable explanations in English


A semantic parser for English

Why hard:

- A large number of possible syntactic parse trees
 - Roles of objects in the sentence hard to understand: subject, object, using, helper, place, time, (“A man opened the door” vs “A key opened the door”)
 - Coreference resolution and friends (“He saw it”)
 - How to represent context/space/time etc?
-

Classic example

Ambiguity



Many frameworks for semantic parsing:

- Abstract Meaning Representation (AMR)
- Universal Conceptual Cognitive Annotation (UCCA)
- Elementary Dependency Structures (EDS)
- Discourse Representation Structures (DRS)
- Universal Decompositional Semantics (UDS)


....

But no good parsers, i.e. suitable for our purposes!


Our approach to semantic parsing: DIY

- We take the output of the Stanza Universal Dependencies Parser (Stanford)
- And then convert the result to 1st order logic enhanced with **confidences and exceptions**.

Stanza uses machine learning on labelled sentences to determine the correct parse.

A blue arrow originates from the text 'Stanza uses machine learning...' and points upwards to the word 'Stanza' in the first bullet point.

The UD framework gives fairly detailed information about the roles of the parts of the sentence.

A blue arrow originates from the text 'The UD framework gives fairly detailed information...' and points upwards to the phrase '1st order logic' in the second bullet point.

We get ca 95% of the Hans adversarial benchmark right: the remaining 5% are Stanza misinterpretations.

Automated reasoner

- Almost all serious work in automated reasoning is either:
 - Research in logic
 - Improving AR for verification (SMT solvers etc)
 - Improving AR for mathematics
- Very little effort has ever gone to automating common sense.

Our commonsense reasoner GK (Graph Knowledge)

GK is built on top of our high-performance classical reasoner GKC

<http://logictools.org/gk>

CADE 2019 paper

CASC prover competition (“the world championship for such systems”) on FOL during FLOC 2022:

First-order Theorems	SnakeFor 1.0	Vampire 4.7	Vampire 4.6	E 3.0	iProver 3.6	CSE_E 1.4	GKC 0.7	Zipperpin 2.1.999	cvc5 1.0	Drodi 3.3.3	CSE 1.5	Prover9 1109a	Goeland 1.0.0	Etableau 0.67
Solved/500	460/500	451/500	448/500	384/500	365/500	361/500	335/500	294/500	271/500	205/500	128/500	123/500	7/500	279/500
Solutions	460 92%	451 90%	448 89%	384 76%	365 73%	361 72%	335 67%	294 58%	271 54%	205 41%	128 25%	123 24%	7 1%	0 0%

Same core

Same core

Our base system GKC for classical FOL

Note: it can be pretty hard to find solutions

- Finding a solution is semidecidable: no guarantee that we can show that a solution does *not* exist.
- Even small problems can be hard. A solvable tiny example from group theory, ca 57 steps in a proof:

$\text{multiply}(A, \text{inverse}(\text{multiply}(B, \text{multiply}(C, \text{multiply}(\text{multiply}(\text{inverse}(C), \text{inverse}(\text{multiply}(D, B))), A)))) = D.$

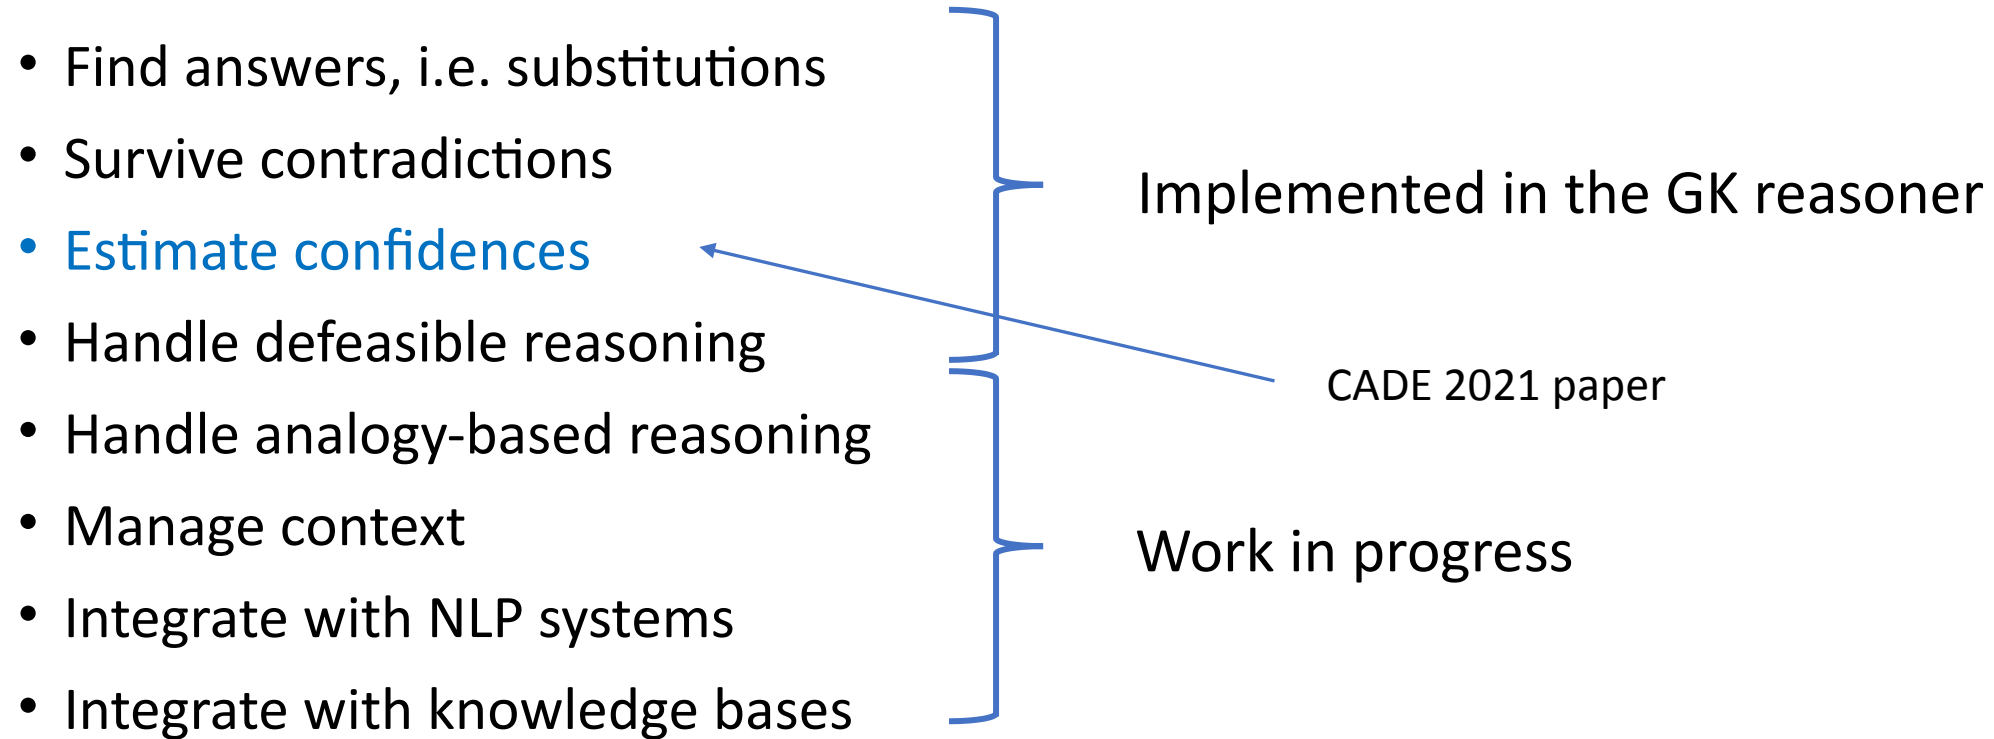
$\text{multiply}(\text{inverse}(a1), a1) \neq \text{multiply}(\text{inverse}(b1), b1).$

Practical context: NLP Q&A / dialogue

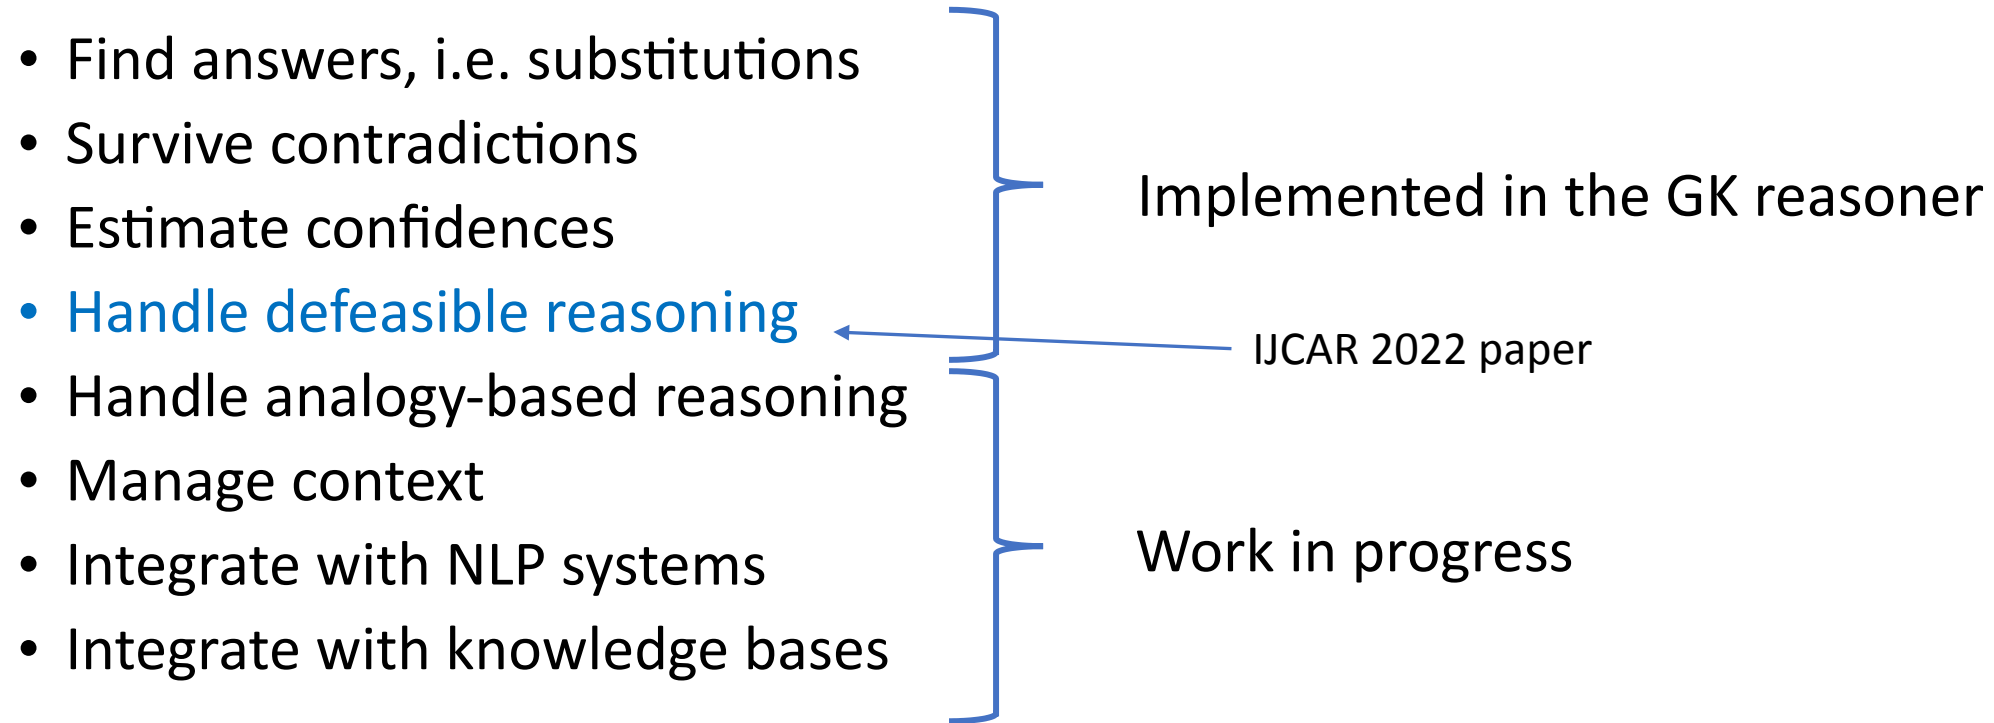
- Suppose we use an AR component as a part of an NLP Q&A / dialogue system.
- Most facts and rules are uncertain: some more, some less.
- We need to estimate our confidence in the results of the derivations.

Using real probabilities does not seem to be realistic for this scenario.

Automated reasoning needs a wealth of capabilities for commonsense A.I. to be usable for a hybrid AR + ML system



Automated reasoning needs a wealth of capabilities for commonsense A.I. to be usable for a hybrid AR + ML system



Recommended listening

- Gary Marcus (KR&R commonsense proponent) & Joshua Bengio (ML proponent)
https://www.facebook.com/watch/live/?v=498403850881660&ref=watch_permalink
read also the related Marcus paper
<https://arxiv.org/ftp/arxiv/papers/1801/1801.00631.pdf> and listen to the Marcus episode on the Lex Fridman podcast <https://lexfridman.com/gary-marcus/>
- David Ferrucci (IBM Watson team leader) on the Lex Fridman podcast
<https://lexfridman.com/david-ferrucci/>
- Tim Rocktäschel (Facebook A.I.) episode
<https://towardsdatascience.com/language-models-symbolic-learning-and-the-road-to-agi-75725985cdf7>
- Reading, not listening: on Hybrid A.I. in image analysis
<https://knowablemagazine.org/article/technology/2020/what-is-neurosymbolic-ai>